

GPU Rendering

This page will help you configure V-Ray for GPU rendering. For help with common questions, please see the [V-Ray GPU Frequently Asked Questions](#) page.

Overview

V-Ray offers two rendering engines - V-Ray and V-Ray GPU.

GPU rendering allows V-Ray to perform the raytracing calculations on the GPU cards installed in the system, rather than the CPU. Since GPUs are specifically designed for massively parallel calculations, they can speed up the rendering process by an order of magnitude. V-Ray GPU uses NVidia CUDA or RTX device(s) to perform the raytracing calculations.

In addition, you can also use CUDA in combination with your CPU device. This is the so called **Hybrid rendering** – when CUDA performs raytracing calculations with the CPU, or simultaneously with both the CPU and GPU devices of your computer.

V-Ray GPU supports a variety of features and even more features are added with time.

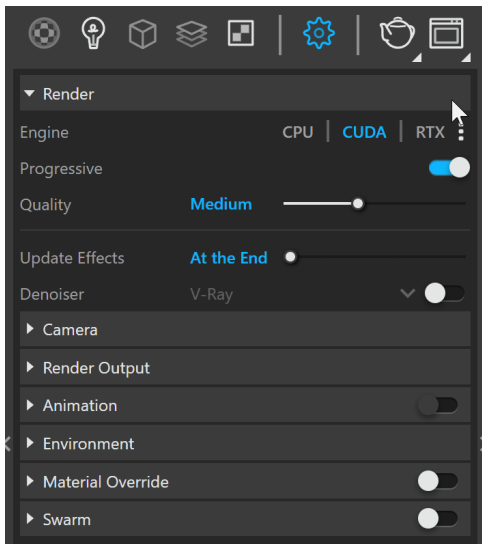
The supported features of V-Ray GPU running on CUDA and RTX are the same.

Although CUDA and RTX share the same user interface as the V-Ray engine, V-Ray GPU differs from the regular V-Ray engine in the way it performs certain calculations. Comparing the results will never come to a one-to-one match, although it may look quite close. Furthermore, it is not the goal for them to be the same.

This is why, it is strongly recommended to not switch between engines in the middle of your project - if you start setting up a scene with the regular V-Ray engine, use it for the entire project. The render settings will only show the available options and your scene will be optimized for GPU rendering.

V-Ray GPU can be used as a production render or in interactive mode to quickly preview scene changes. It also supports both the Progressive and Bucket Image Samplers.

To enable GPU rendering, from V-Ray Asset Editor Settings tab Render rollout, select **CUDA** or **RTX** engine. You can use it with both Progressive and Bucket Sampler types.



Why V-Ray GPU?

V-Ray GPU has a number of advantages:

- Graphics cards often outperform CPU devices when it comes to processing parallel tasks, such as raytracing;
- A single machine can host multiple GPU devices but in most cases only a single CPU;
- Hybrid rendering allows maximum hardware utilization by harnessing the computing power of both GPU and CPU devices;
- V-Ray GPU has all the required features to be production ready;

V-Ray CPU advantages:

- Allows rendering complex scenes requiring a substantial amount of memory;
- Supports all V-Ray features;
- Does not pre-occupy single-GPU machines during rendering

Supported Hardware and Drivers

V-Ray GPU offers rendering with NVIDIA CUDA and RTX platforms.

Below is a list of tests we have performed and the results:

- **NVIDIA** – The CUDA engine is supported only in 64-bit builds of V-Ray for Maxwell and later NVIDIA cards. Rendering on multiple GPUs is also supported. See [here](#) if your card has the minimum required compute capability.
- **NVIDIA RTX** – Choosing RTX GPU mode works with RTX cards.
- **Hybrid Rendering (running CUDA on GPU and CPU)**: Starting with V-Ray 3.6, V-Ray GPU CUDA rendering can be performed on CPUs and NVIDIA GPUs at the same time. Using the Select Devices for V-Ray GPU Rendering you can enable your CPUs as CUDA devices and allow the CUDA code to combine your CPUs and GPUs to utilize all available resources.

Rendering on multiple GPUs is supported, and by default, Interactive rendering for GPU uses all available GPU devices. See the sections below to learn how to choose devices on which to run IR GPU.

If V-Ray GPU cannot find a supported CUDA device on the system, it silently falls back to CPU code. If V-Ray GPU cannot find a supported RTX device on the system, the process stops.

To see if the V-Ray render server is really rendering on the GPU, check out its console output.

When rendering, the output window provides information about your current driver version. If the version is not compatible with V-Ray GPU, you will find a message with the recommended version.

o use NVLINK on supported hardware, NVLINK devices **must** be set to TCC mode. This is recommended for Pascal, Volta and Turing-based Quadro models. For GeForce RTX cards, a SLI setup is sufficient. Also note that to prevent performance loss, not all data is shared between devices.

For more information, see the OptiX and NVLink FAQ page.

The RTX engine type only works with NVidia RTX cards.

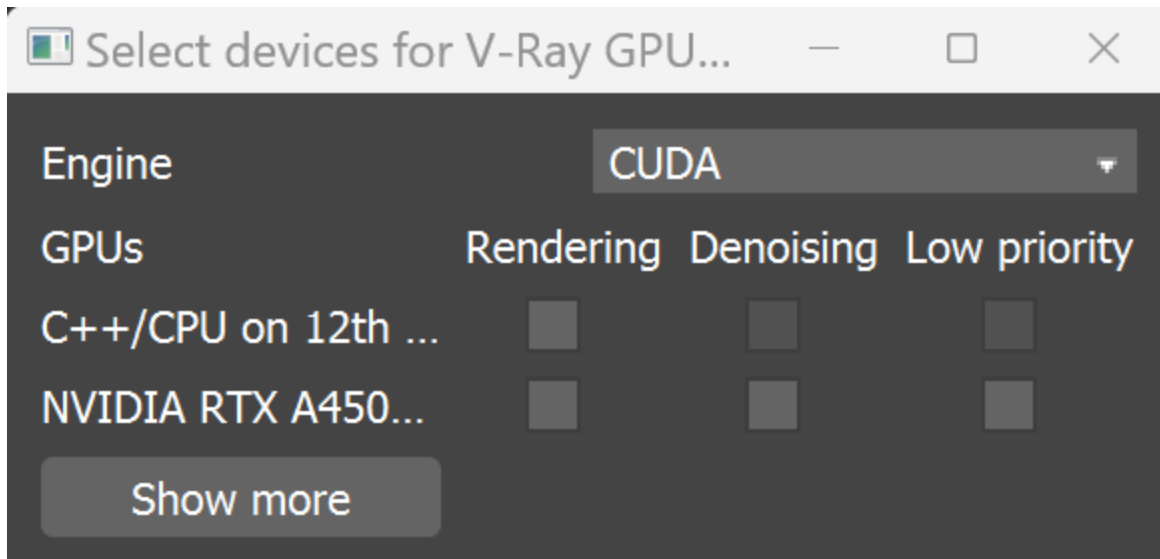
Choosing Which Devices to Use for Rendering

You may not want to use all available GPU devices for rendering, especially if you have multiple GPUs and you want to leave one of them free for working on the user interface or you may want to combine your CPU and GPU together (see the [Hybrid Rendering section](#) below).

There are two ways to do this: by using the supplied standalone device selection tool, accessible from both the V-Ray Tools menu and the Start menu search, or from the Asset Editor's GPU device dropdown list.


When selecting devices from the standalone tool, that will set the default setup for the work station. After changing this option, you need to restart Rhino and/or any Standalone DR servers (if any are running) for the changes to take effect. If the Standalone DR server is running as a service, you may need to restart it.

When choosing the devices from the Asset Editor - it is done only for the current session in Rhino.



||V-Ray||> **Tools** > **External** > **GPU Device Selection**

Or find the `vray_gpu_device_select.exe` file here [C:\Program Files\Chaos Group\V-Ray\V-Ray for Rhinoceros\vraysdk\bin.](#)



||V-Ray Asset Editor||> **Settings** > **Render** > **GPU** device list dropdown

If you see your CPU listed twice, choose the option with "C++/CPU " in the name.

Hybrid Rendering with CPUs and the CUDA Engine

Starting in 3.60, V-Ray GPU can perform hybrid rendering with the CUDA engine utilizing both the CPU and NVIDIA GPUs. V-Ray can now execute the CUDA source on the CPU, as though the CPU was another CUDA device. To enable the hybrid rendering mode, simply enable the **C++/CPU** device from the list of CUDA devices.

The hybrid rendering mode does not require any special drivers. Furthermore, you can use the CPU as a CUDA device even if you don't have an NVIDIA GPU and/or NVIDIA drivers installed. Meaning, this mode can be used on computers that don't even have GPUs. The hybrid render engine running on a CPU [supports the same features](#) as the regular V-Ray GPU CUDA engine.

More info about the nature of Hybrid rendering is available at the blog post [Understanding V-Ray Hybrid rendering](#).

Notes

Memory usage can be viewed from the [Stats section](#) of the new V-Ray Frame Buffer.

References

- The NVIDIA CUDA developer zone: <https://developer.nvidia.com/category/zone/cuda-zone>
- For more information on hardware recommendations, see the [Hardware Recommendations](#) article.